# Measuring the effects of DNSSEC deployment on query load

Jelte Jansen[*]
NLnet Labs

NLnet Labs document 2006-002 May 1, 2006

**Abstract**

Ripe NCC recently started signing the zones on their DNS servers. This document presents a few measurements of the effects (if any) on the behaviour of the resolvers sending queries to the Ripe nameservers. We have looked at the rate of queries with the DO bit ('use DNSSEC') set to 1, compared to those with the DO bit set to 0. We have also looked at the number of DNS responses that were truncated.

## Contents

[*] jelte@nlnetlabs

File information:Id: ripe-signed-zones.tex 35 2006-05-01 15:16:49Z jeltejan

# 1 Introduction

DNS Security Extensions (DNSSEC) [1, 3, 2] add integrity checking to the domain name system. DNS Resource Records (RR) are signed using private keys. The signatures are published in the DNS as RRSIG resource records. The public keys that are needed to validate the signatures are published as DNSKEY resource records.

The introduction of these records might cause resolvers to make a lot of extra queries. Especially if those resolvers use broken implementations that incorrectly signal that they can handle DNSSEC data. What these resolvers do if they get DNSKEY and RRSIG and NSEC records is unknown. It could very well be that there is an explosion of queries from those resolvers, since they might not be able to parse the result they get, and they will always get the same results. A lot of resolvers tend to just resend their query if they do not get their answer.

Another possible problem is that packets are dropped by firewalls that filter on the protocol level. If they do not recognize the DNSSEC resource records, they might drop the packet.

In this paper, we address the following question: *What is the immediate effect of signing zones, measured by the number of queries sent for those zones?*

We performed the analysis on actual tcpdump data of the nameservers. From each day in the measuring period, a tcpdump is made for a fixed time period (starting at 15:00:01 and ending at 15:10:00 every day). These traces are examined with ldns-dpa, a dns packet analyzer. [1]

[1] ldns-dpa is one of the example tools of the NLnet Labs DNS library libdns. The next public version will include this tool.

## 2 Queries with and without the DO bit

In this chapter, we show a comparison of the number of packets with the DO bit [2] set to 0 and those with the DO bit set to 1. We do this for the total number of queries and for all queries for a number of zones. The blue line marks the date on which a zone was signed. This date is determined by the first encounter of RRSIG Resource records in DNS answers.

Only UDP packets are evaluated. See chapter 3 for a measurement of truncated UDP packets.

Since the data of only ten minutes each day is used, the differences between two successive days can be significant.

[2] DO = "DNSSEC OK", signals that the resolving software wants to receive DNSSEC data. See [4]

## 2.1 Data

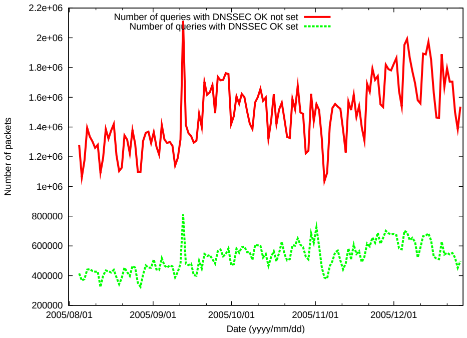

*Figure 1: All zones*

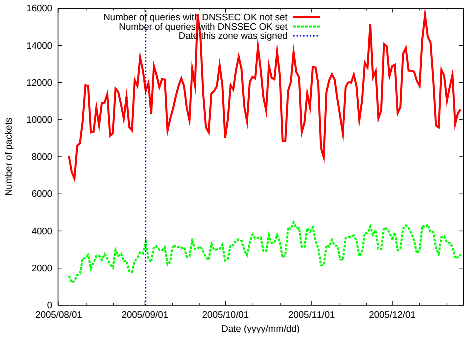

*Figure 2: ripe.net*

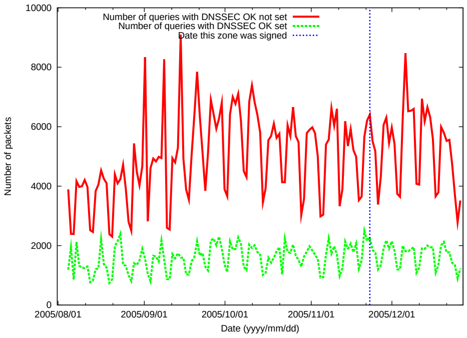

*Figure 3: 145.in-addr.arpa*

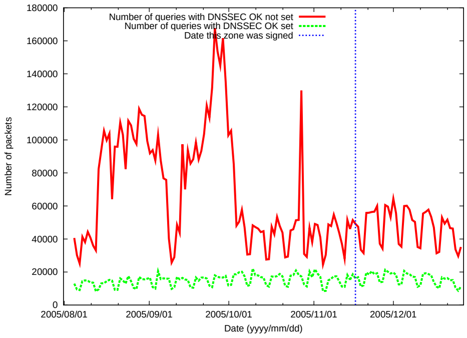

*Figure 4: 193.in-addr.arpa*

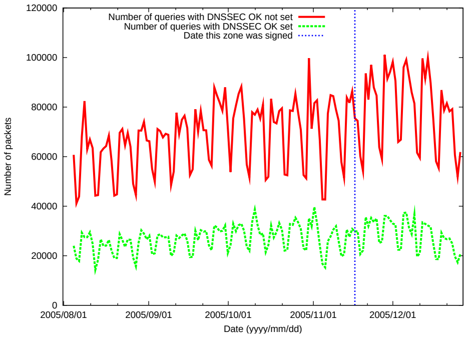

*Figure 5: 195.in-addr.arpa*

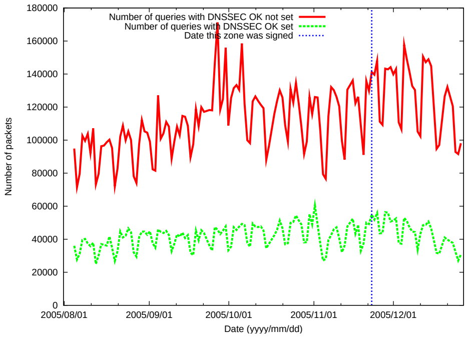

*Figure 6: 212.in-addr.arpa*

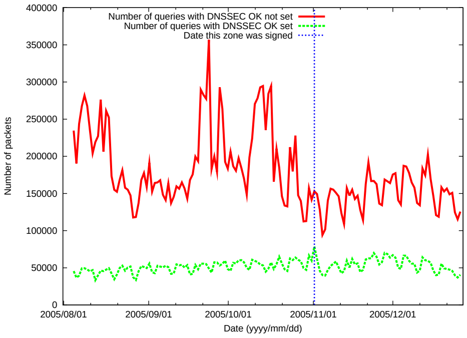

Figure 7: 213.in-addr.arpa

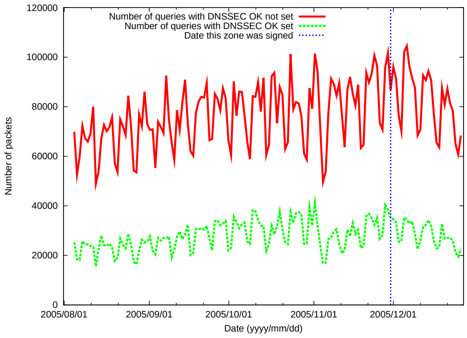

Figure 8: 217.in-addr.arpa

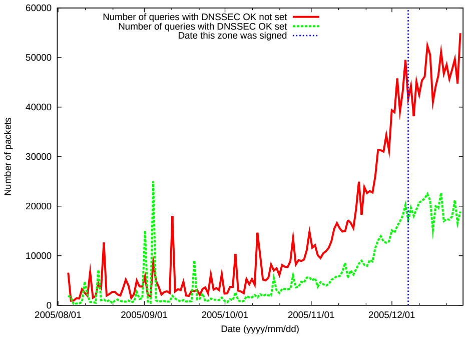

Figure 9: 88.in-addr.arpa

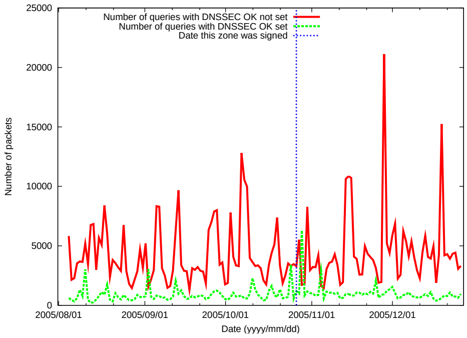

Figure 10: 89.in-addr.arpa

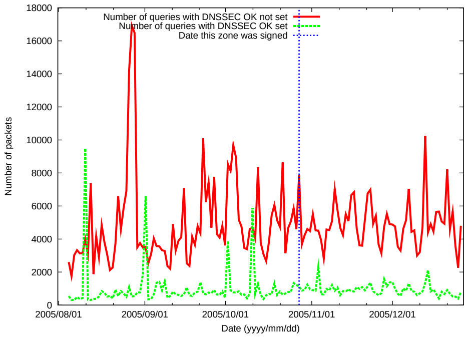

Figure 11: 90.in-addr.arpa

## 2.2 Discussion

Overall, there is a relatively slight increase in query load, but not in the number of DO=1 queries relative to the number those with DO=0.

The 88.in-addr.arpa zone does show a significant increase in queries, but that increase started before the zone was signed. The increase is probably caused by deployment of networks within the zone.

# 3 Truncated answer packets

Since the introduction of DNSSEC in these zones, there has been an increase in truncated packets (with the TC bit [3] set to 1) and hence in the number of TCP queries. The following diagrams show the number of truncated packets for each zone per day.

The red line shows the number of packets with the TC bit set to 1. Since the number of truncated packets is close to zero compared to the number of packets with the TC bit set to 0, we will show those in separate diagrams.

Where the red line is not present, no packets with the TC bit set to 1 were found.

The green line marks the day the zone was signed.

[3] TrunCation, the length of the message was greater that permitted on this transmission channel, see [5]

## 3.1 Data

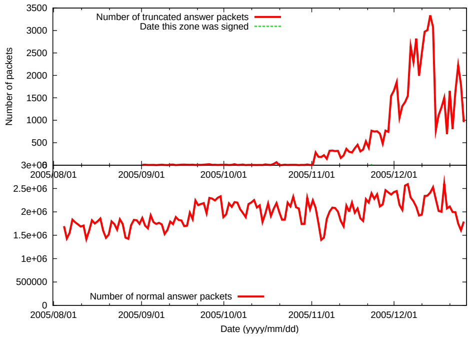

*Figure 12: All zones*

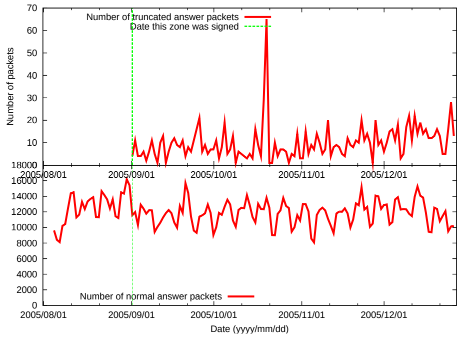

*Figure 13: ripe.net*

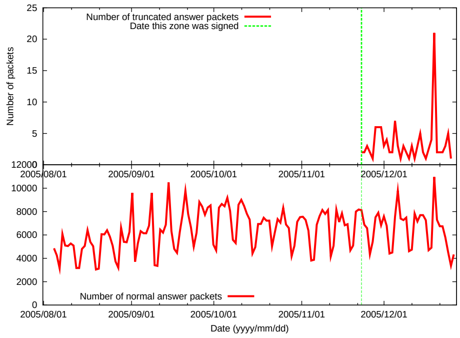

*Figure 14: 145.in-addr.arpa*

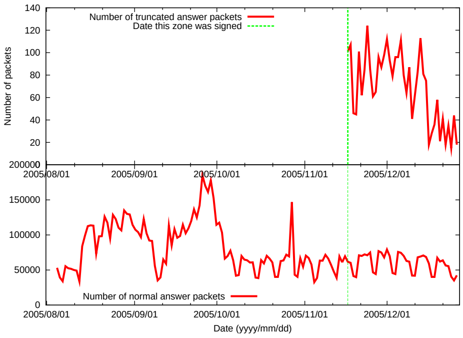

*Figure 15: 193.in-addr.arpa*

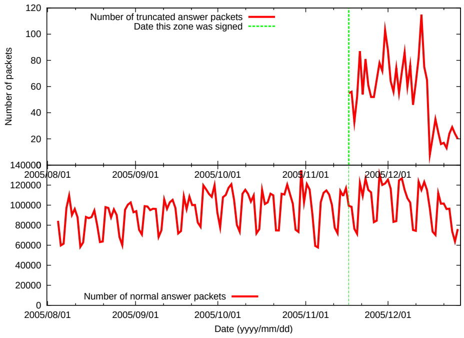

*Figure 16: 195.in-addr.arpa*

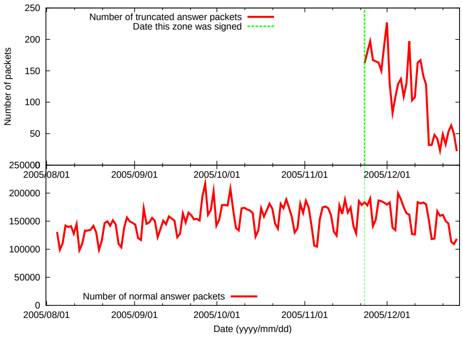

*Figure 17: 212.in-addr.arpa*

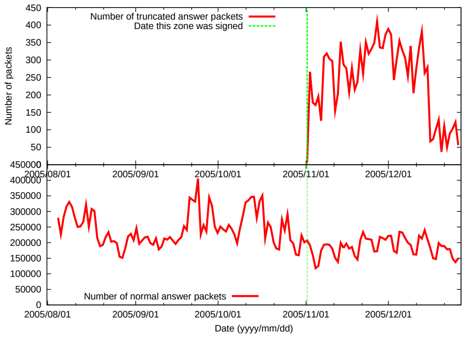

*Figure 18: 213.in-addr.arpa*

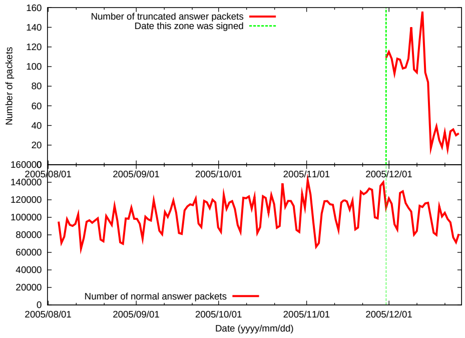

*Figure 19: 217.in-addr.arpa*

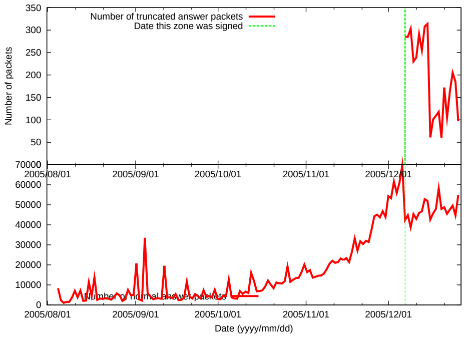

*Figure 20: 88.in-addr.arpa*

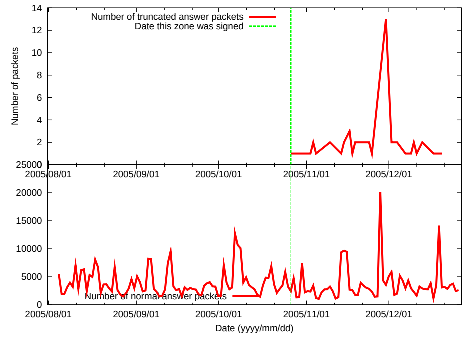

*Figure 21: 89.in-addr.arpa*

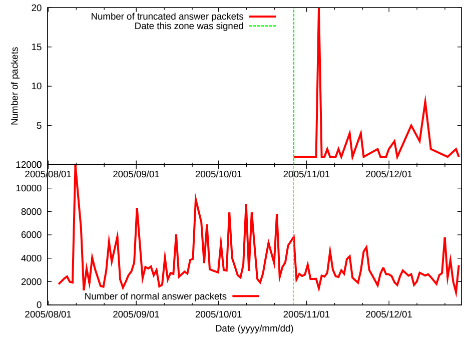

*Figure 22: 90.in-addr.arpa*

## 3.2 Discussion

As can be seen in the diagrams, as soon as zones are signed, truncated packets start to appear.

Apart from a few exceptions, all truncated answers were of the type name error, and without exception, all queries that resulted in runcated answers had the DO bit set to 1, and their UDP size length set to 512 bytes. 512 bytes is in these cases not enough to put in the necessary authority data and the signatures for that data. Most of the inspected truncated answers are missing the signature for the NSEC data.

We have counted 92 individual hosts that sent queries with the DO bit set and and EDNS packet size of 512. We have been able to fingerprint 55 of these hosts, all were running BIND 9.2.3rc1 - 9.4.0a0. We were not able to fingerprint the other 37.

These versions of Bind have a default EDNS0 buffer size setting of either 2048 or 4096, and versions 9.3.0 and later support DNSSECbis, so it appears that the edns-udp-size has been manually set to 512 bytes.

Although we cannot be sure if the fingerprinting is accurate, since we do not have actual deployment data, these results seem feasible. From the BIND9 manual:

> **edns-udp-size**
>
> sets the advertised EDNS UDP buffer size. Valid values are 512 to 4096 (values outside this range will be silently adjusted). The default value is 4096. The usual reason for setting edns-udp-size to a non default value it to get UDP answers to pass through broken firewalls that block fragmented packets and/or block UDP packets that are greater than 512 bytes.

This text could cause operators to set the value to 512 by default.

We have inspected some of these cases more closely, and upon receiving a truncated answer, all resolvers we have looked at responded immediately with a query over TCP. This indicates that the existence of DNSSEC data in answers in these cases does not cause the packets to be dropped by intermediate machines such as firewalls. Keep in mind that this was only a sample, and we have little information about deployed software at the client side.

A few resolvers were seen that did not immediately query again after the TCP session, but did perform the same query a while later. The rate of requerying varied from a few seconds to a few minutes. Due to the short timeframe of our data, we cannot say anything more about the cause of this behaviour.

# 4 Conclusion

At this moment, there is no immediate noticeable increase in the total number of queries that are sent, either with or without the DO bit set.

There is an increase in truncated answers, consisting almost completely of name error answers with DNSSEC data. The number of truncated answers is minimal compared to the total number of answers.

Examination of the queries that resulted in these answers shows that these queries have the DO bit set, but an edns-udp-size that was set to 512. This data suggests that the software is working correctly but the configuration causes packets to be truncated.

The data does not show signs of nonlinear effects caused by packets that are dropped between authoritative servers and recursive servers because DNSSEC was enabled on the server.

## References

[1] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *DNS Security Introduction and Requirements.* RFC 4033 (Proposed Standard), March 2005. http://www.ietf.org/rfc/rfc4033.txt.

[2] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Protocol Modifications for the DNS Security Extensions.* RFC 4035 (Proposed Standard), March 2005. http://www.ietf.org/rfc/rfc4035.txt.

[3] R. Arends, R. Austein, M. Larson, D. Massey, and S. Rose. *Resource Records for the DNS Security Extensions.* RFC 4034 (Proposed Standard), March 2005. http://www.ietf.org/rfc/rfc4034.txt.

[4] D. Conrad. *Indicating Resolver Support of DNSSEC.* RFC 3225 (Proposed Standard), December 2001. http://www.ietf.org/rfc/rfc3225.txt, (Updated by RFCs 4033, 4034, 4035).

[5] P.V. Mockapetris. *Domain names - implementation and specification.* RFC 1035 (Standard), November 1987. http://www.ietf.org/rfc/rfc1035.txt, (Updated by RFCs 1101, 1183, 1348, 1876, 1982, 1995, 1996, 2065, 2136, 2181, 2137, 2308, 2535, 2845, 3425, 3658, 4033, 4034, 4035).